

Causal Models and Exploratory Analysis in Heterogeneous Information Fusion for Detecting Potential Terrorists

Paul K. Davis, David Manheim, Walter L. Perry, and John S. Hollywood

RAND National Security Research Division

WR-1124

November 2015

RAND working papers are intended to share researchers' latest findings and to solicit informal peer review. They have been approved for circulation by (name of RAND research unit) but have not been formally edited or peer reviewed. Unless otherwise indicated, working papers can be quoted and cited without permission of the author, provided the source is clearly referred to as a working paper. RAND's publications do not necessarily reflect the opinions of its research clients and sponsors. RAND® is a registered trademark.



Preface

This Working Paper was prepared for the third annual *Workshop on Decisionmaking Under Deep Uncertainty*, to be held in Delft, the Netherlands, November 3-5, 2015. The Working Paper is based on a RAND report that will be published shortly. Informal comments on the Working Paper are welcome and should be addressed to me at pdavis@rand.org.

This research was conducted within the International Security and Defense Policy Center of the RAND National Security Research Division (NSRD). NSRD conducts research and analysis on defense and national security topics for the U.S. and allied defense, foreign policy, homeland security, and intelligence communities and foundations and other non-governmental organizations that support defense and national security analysis. For more information on the International Security and Defense Policy Center, see <http://www.rand.org/nsrd/ndri/centers>.

Table of Contents

Preface.....	ii
Abstract.....	1
1 Introduction.....	1
2 Top-Level Analytical Architecture.....	2
3 Representing Heterogeneous Information	4
3.1 Types of Information	4
3.2 Representing Uncertainty.....	4
4 Causal Social-Science Models in Counterterrorism	6
5 A Mixed-Methods Battery of Fusion Methods.....	11
6 Data.....	12
7 Designing and Implementing a Platform for Exploratory Analysis	13
8 Illustrative Results and Conclusions.....	17
8.1 Results.....	17
8.2 Conclusions.....	18
Acknowledgments.....	19
References.....	19
Author Biographies.....	20

Abstract

We describe research fusing heterogeneous information in an effort eventually to detect terrorists, reduce false alarms, and exonerate those falsely identified. The specific research is more humble, using synthetic data and first versions of fusion methods. Both the information and the fusion methods are subject to deep uncertainty. The information may also be fragmentary, indirect, soft, conflicting, and even deceptive. We developed a research prototype of an analyst-centric fusion platform. This uses (1) causal computational models rooted in social science to relate observable information about individuals to an estimate of the threat that the individual poses and (2) a battery of different methods to fuse across information reports. We account for uncertainties about the causal model, the information, and the fusion methods. We address structural and parametric uncertainties, including uncertainties about the uncertainties, at different levels of detail. We use a combination of (1) probabilistic and parametric methods, (2) alternative models, and (3) alternative fusion methods that include nonlinear algebraic combination, Bayesian inference, and an entropy-maximizing approach. This paper focuses primarily on dealing with deep uncertainty in multiple dimensions.

1 Introduction

This paper has substantive and methodological purposes. Substantively, we describe basic research on using information fusion to assess whether individuals pose a terrorism threat. Methodologically, we illustrate how causal social-science models can be used as part of doing so. We also describe our analytical architecture for exploratory analysis under multi-dimensional deep uncertainty.

Earlier work reviewed technologies for using behavioral observations as part of threat detection (Davis, Perry, Brown, Jeung, et al, 2013). A conclusion was that some of the technologies were promising, but magic bullets were unlikely and future success would probably depend on greatly improving methods for information fusion. Unlike fusion in many domains, however, the information in question is mixed in character and often both complex and soft—i.e., qualitative, subjective, fuzzy, or ambiguous. It is likely also to be uncertain, conflicting, and sometimes deliberately deceptive. Further, the underlying phenomena are more poorly understood than in physics problems. All of this motivated our current research, the report for which is in the publications process (Davis, et al., forthcoming).

We sought to build a prototype system that would fuse reports on an individual to be assessed as a possible terrorist threat. It should, we concluded:

- (1) Address matters probabilistically (how likely is it that the individual poses some *degree* of threat?);

- (2) Employ *causal* social-science models as part of the fusion;
- (3) Use a mixed-methods approach to span the range of reasonable causal relationships and fusion methods;
- (4) Have competitive streams of analysis because results depend on human imagination and judgment;
- (5) Design for routine exploratory analysis under deep uncertainty; and
- (6) Embody the system in a comprehensible analyst-centric computer platform that would facilitate review, debate, and eventual sharing and re-use.

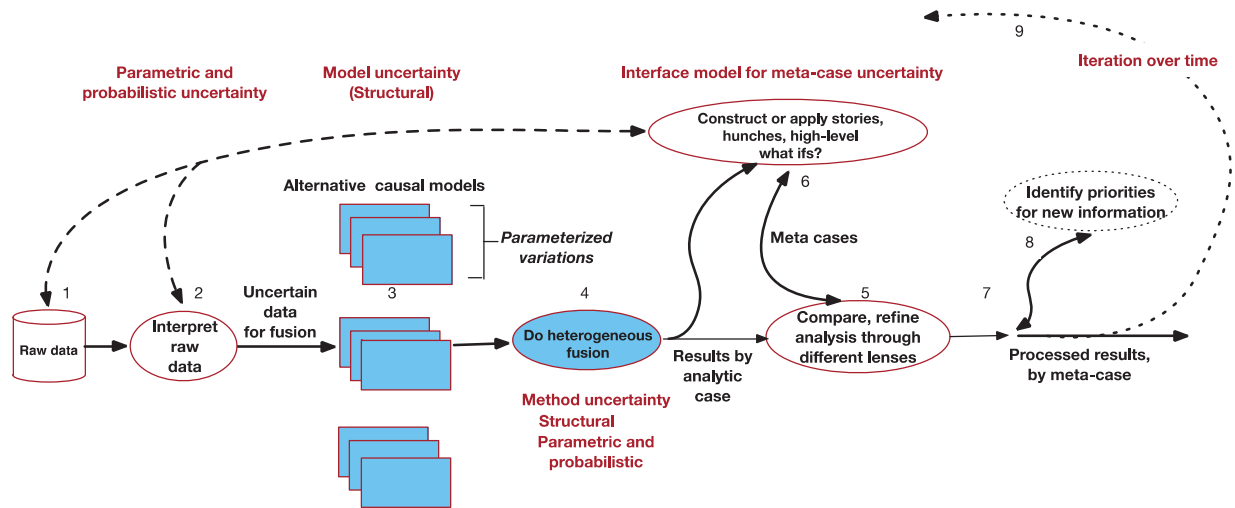
We focused on terrorism, but the research should be relevant to law enforcement, intelligence, and other domains. Our approach is very different from machine-driven mining of information from “Big Data.”

Subsequent sections describe our analytical architecture, representation of heterogeneous information and related uncertainties, causal models, mix of fusion methods, creation of synthetic data for testing, and implementation challenges. We end with illustrative results and conclusions.

2 Top-Level Analytical Architecture


Figure 1 shows the conceptual analytical architecture, which embodies the characteristics identified above. When assessing an individual for the degree to which he poses a threat of terrorism, information derives from various sources as shown on the far left (1). It must be “interpreted” to create the inputs needed for subsequent processing (2). The interpreted data then enters the several separate streams of analysis, with different causal models and variants (3). Information result is then fused (4), drawing on an entire suite of methods and tools. The results, by method, are passed along for comparisons and for refinement of analysis with case-dependent tuning (5)-(6). This leads to results (7), which are also starting points for iteration with new information or the revisiting of old raw data (9). In the future, the analysis will also be used to identify priorities for obtaining new information to sharpen analysis (8).

Figure 1. Top-Level System Vision



An important element (6) is labeled “interface model.” This is where the analyst can change elements of the fusion process, perhaps to reflect different interpretations of the original raw data or to generate displays responsive to high-level questions that have to be mapped into the inputs of the fusion system. To illustrate, suppose that several “stories” are floating around in peoples’ heads about how to understand the information coming in. These might be akin to what a detective novel might call different theories of the case. They may also stem from intuition or biases, depending on perspective. If they are made explicit, then an “interface model” can be written that says, e.g., “If we want to see things through the lens of story 1, then ...” As indicated in Figure 2, this may affect some of the priors for Bayesian analysis, the form of some input probability distributions, the way in which missing information is dealt with, and parameter settings. It can also cause the analyst to look harder at some fusion methods than others. Much of this corresponds to nontrivial 1:n mappings requiring knowledge, imagination, and artistry. We should never imagine that such information fusion is a mechanical and objective process.

Figure 2. How Different “Stories” Can Affect Fusion



3 Representing Heterogeneous Information

3.1 Types of Information

Heterogeneous information comes from different sources (e.g., detection devices, digital records, and human sources) and also varies in character. It can be complex and “soft”—i.e., qualitative, subjective, fuzzy, or ambiguous—and also contradictory or even deceptive. Human sources sometimes lie, sometimes to curry favor and sometimes with malicious intent. Sensor data and archival digital records can simply be wrong. Information can also be complex, as in “He is either a committed terrorist or a good-hearted naïf in bad company where he acts as he is expected to act in that group. I’m not sure which” (resulting in a bimodal probability distribution). Another example would be “I’m not sure whether his motivation for the cause is high or very high, but it’s up there.”

The problems of dealing with complex information have been discussed for decades, some of it debating about Bayesian, Dempster-Shafer, Possibility Theory, Dezert-Smarandache Theory, and other methods; some of it in connection with “fuzzy mathematics.” We constructed a lengthy list of complex information types to challenge ourselves. We found that we could represent all the classes of complex information with the methods that we adopted. These included (1) going beyond binary thinking; (2) using probability distributions for report data; (3) using causal models; (4) using fusion methods that allow inequalities and either promote convergence or preserve distinctions, depending on context; (5) representing qualitative variables on a common 0 to 10 scale (or its discrete version $\{1.3.5.7.9\}$); and (6) characterizing the credibility and salience of information elements. In principle, these can also be achieved with a more complex Bayesian-network approach.

3.2 Representing Uncertainty

Representing uncertainty was a major issue. We had to deal with model uncertainty (structural uncertainty) and parametric uncertainty. These, however, appeared in numerous places as indicated by Table 1. We had various causal models, fusion methods, and ways to combine them (a meta-model of fusion). Even when dealing with structural uncertainty, we also had parametric choices. And, as discussed in Figures 1 and 2, we needed some models of our data to deal with higher-level issues and questions.

We used a variety of methods to deal with uncertainty (Table 2) and were careful to maintain distinctions between source data and methods used. Multiple methods were applied for each uncertainty type: we varied parameter values; we employed probability distributions, both for representing gaps in knowledge and effects of random processes; and we used combinations. We used alternative causal and fusion models. And, finally, we used interface models. Other authors have encountered the same issues and made some of the same distinctions (Walker, Lempert, and Kwakkel; Bankes, Lempert, and Popper, 2005), but it did appear that we had more than usual

complexity. This was because we had no constraints—no sponsor insisting that we worry about only some uncertainties. We also wanted to make the various choices easy for an analyst to adjust, which can be difficult with other methods for exploratory analysis.

Table 1. Types of Uncertainty or Disagreement for Different Elements of Platform

Type of Uncertainty	Causal model of phenomenon	Causal-model data	Model of causal-model data	Fusion model	Data for fusion model (tuning parameters)
Structural	•		•	•	
Parametric	•	•		•	•

As indicated in Table 2, we distinguished between uncertainties due to random processes and uncertainties due to lack of knowledge. Both can be represented by probability distributions, but they are different and the differences matter. This was a continuing source of confusion because many researchers refer to both classes in terms of random variables without noting subtleties involved (e.g., the need to use mixture distributions for some of the work).

Table 2. Methods for Dealing with Uncertainty

Mechanism	Causal model of phenomenon	Causal-model data	Model of causal-model data	Fusion model	Data for fusion model
Deterministic parameter variation	•			•	
Chunky variations (e.g., different data tables)	•••	•••			•••
Probability distributions for knowledge		•••			
Probability distributions for random processes	•				
Alternative models	•••			•••	
Interface model		•••	•••		

4 Causal Social-Science Models in Counterterrorism

The causes for an individual becoming a terrorist are complex and only partially understood. To compound this uncertainty, if an individual is being evaluated for the threat of terrorism he poses, the evidence is likely to be about observed or inferred attributes that might suggest that he represents some degree of threat (to include zero). For example, he might have military expertise with explosives and a high degree of interest in radical Islamic philosophy. To use available information, we needed a causal model rooted in social science to represent causal factors and their effects.

An earlier review of scholarly social science bearing on terrorism synthesized in terms of easy-to-understand “factor trees” identifying the factors contributing in a causal way to terrorism and public support thereof (Davis and Cragin 2009). Factor trees are “thinking models” to structure reasoning and discussion. Case studies were then conducted to test or “validate” a particular factor tree, that for public support of terrorism. This factor tree (Figure 3) held up well as a general qualitative theory with myriad context-specific specializations (Davis, Larson, et al.,

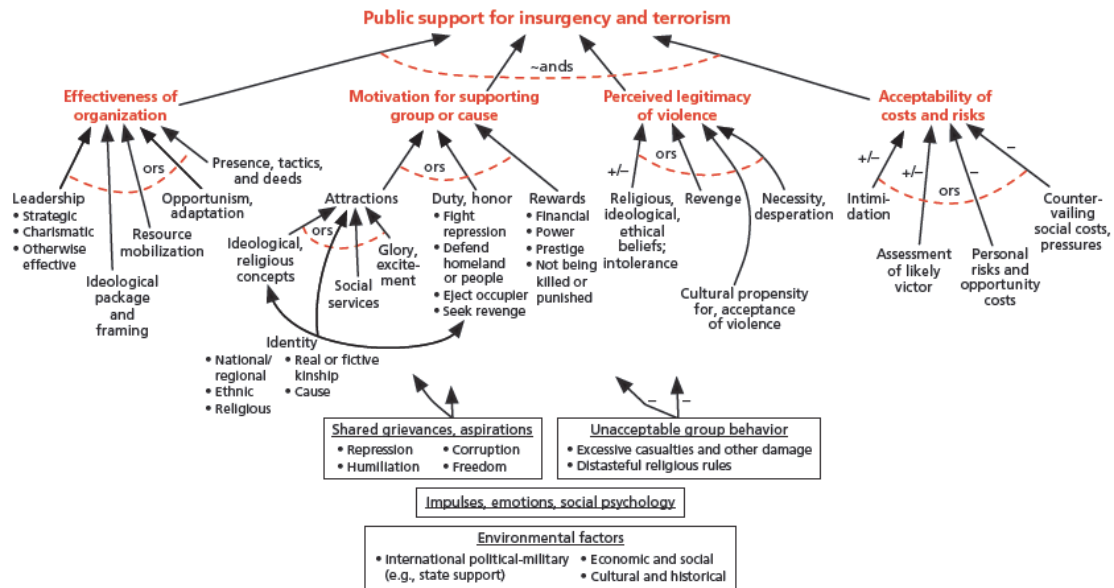
2012). The study discussed what “validation” does and does not mean in such work, a matter of interest to all concerned with deep uncertainty. It is one thing to validate that a model has captured the right factors; it is quite another to predict combined consequences.

The first point from Figure 3 is that a given higher level factor such as motivation can have many sources: religious zealotry is *one*, but so also the sources may be a sense of identity, a desire for glory and excitement, or a sense of duty. Recognizing such multiple causes changes discussion from arguing about which single cause is correct to something more realistic. The second point is that factor tree use “~ands” to indicate that, to a first approximation, where *all* of the contributing factors need to be present if the factor to which they contribute is to be significant. They use “ors” to indicate where any or all of contributing factors may be sufficient to create the higher-level factor (i.e., where the contributing factors are substitutable).

Figure 3 is a multi-resolution model (MRM) (Davis 2003): one can specify inputs at the top level, at the level with four main factors, at the next more detailed level, etc. Making such relationships explicit is very helpful conceptually and also for empirical analysis. In empirical work, whether using real-world or model-generated data, higher level factors are good theory-informed abstractions. Those—not simple-hypothesis variables—should often be the core of multivariate regression.

Factor trees are qualitative, but it proved possible subsequently to build a computational model based on the factor tree of Figure 3. This was an unusual computational model, however. *It was developed exclusively for exploratory analysis under deep uncertainty, rather than making best-estimate predictions* (Davis and O'Mahony 2013).

Figure 3: A Factor Tree for Public Support of Insurgency and Terrorism

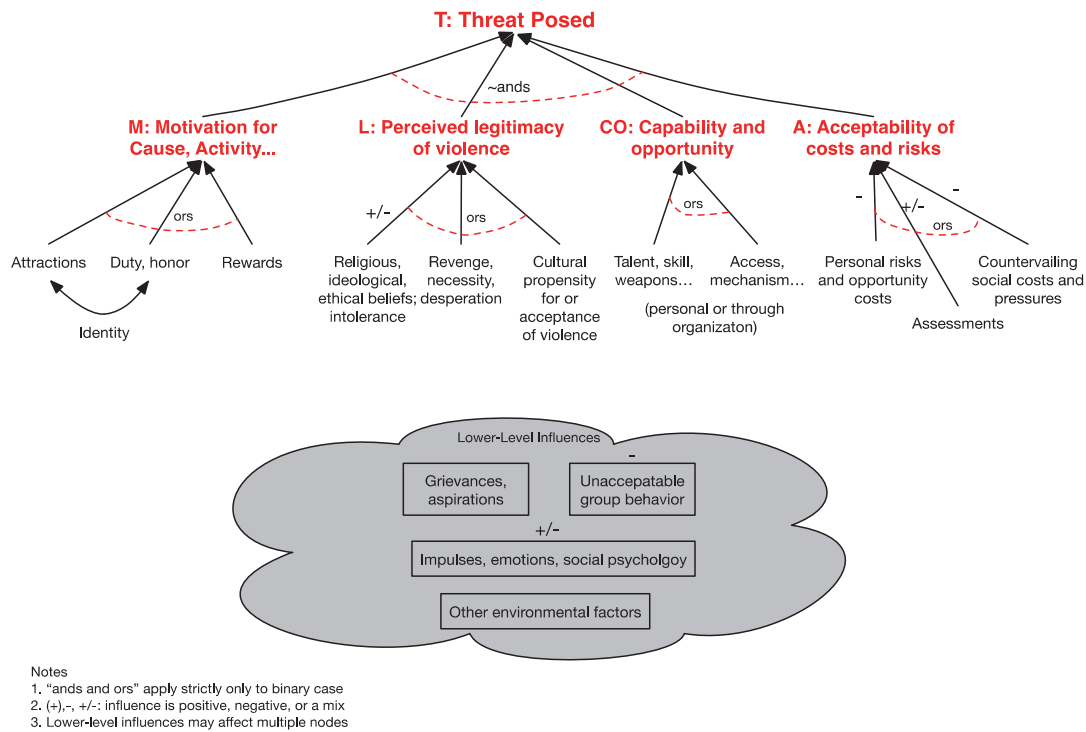


NOTES: Applies at a snapshot in time. Current factor values can affect future values of some or all other factors.

RAND TR1220-S.1

For the present work we extended the Davis-O'Mahony methods to develop the Propensity for Terrorism (PFT) model, a truncated version of which is shown in Figure 4. It depicts the factors influencing an individual's propensity to commit terrorism, and, thus, the threat posed by the individual. Although not separately validated by social-science research, it builds on the earlier work and seemed as a reasonably credible example. We used only the top layer that asserts that the threat T posed by an individual is a function of that individual's motivation (M), perception of terrorism's legitimacy (L), capability-opportunity (CO), and acceptability of costs (A). These inputs can be informed by observations or inferences about the lower-level factors in the tree, or by interpreting other data that might not connect straightforwardly to the causal model. This would include using statistical information to inform base rates.

Figure 4. A Truncated Factor Tree of Propensity for Terrorism



We defined *M*, *L*, *CO*, and *A* so as to make them essentially independent, both logically and probabilistically. This depends, however, on the model being used with data that has been properly interpreted. Thus, we allow, by exception, for explicit correlations to be specified. For example, if a given human source does not understand the model's difference between motivation for a cause and a sense of legitimacy in using terrorism tactics, then he may report *M* and *L* as the same. Thus, his inputs on *M* and *L* would not be independent.

The PFT model is deliberately not that of a rational actor doing cost-benefit calculations. Real individuals are affected by emotions, cognitive biases, and other non-rational considerations. Further, they may not even have stable utility functions, instead "discovering" their values as matters develop, rather as with "wicked problems" in the policy domain. Such issues are discussed, with pointers to the literature, in a recent National Academy report about deterrence (National Research Council 2014, 35ff) to which one of us (Davis) contributed.

Briefly, some technical elements of the PFT model are as follows.

We define the factors (variables) on an interval scale of 0 to 10, often using the discretized scale of 1, 3, 5, 7, 9 with equally spaced values corresponding to very low, low, medium, high, very high.

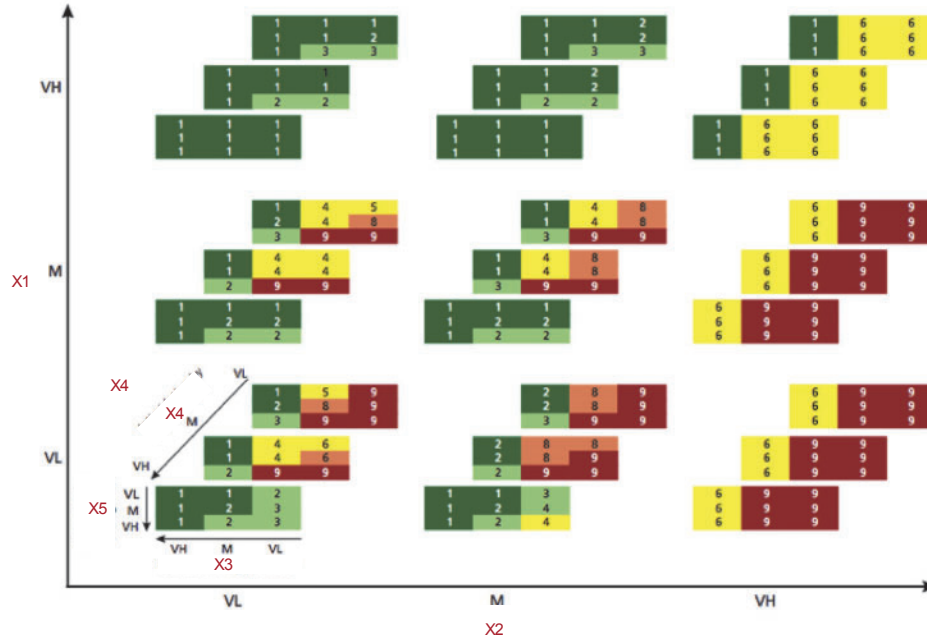
Although we don't know the potentially complex actual functions describing the combined effects at each node of a factor tree, we find that much can be accomplished with a

combination of two building-block functional forms that more or less bound ways to represent nonlinear effects. We call them Thresholded Linear Weighted Sums (TLWS) and Primary Factors (PF).

The TLWS method is akin to ordinary linear weighted sums, but it sets the result to 0 unless *each* of the contributing variables exceeds its threshold. This addresses a number of practical nonlinearities in a simple way. For example, if someone will very likely not become a terrorist unless he is willing to accept associated costs and risks, then there is some threshold level constituting “enough” acceptance.

The Primary Factors (PF) method states that the result is determined by the largest of the contributing factors, or that plus a minor upward adjustment reflecting the next-biggest. For example, an individual might have numerous sources of motivation, but the strength of his motivation might be largely dictated by the biggest single one (e.g., extremist religion or a love of danger and violence).

We developed the model (and the larger platform) with the *Analytica*® modeling platform. It uses visual modeling that is largely understandable to people who are not “real” programmers. Also, its modeling paradigm is closely related to the natural mathematics of the problem: it exploits array mathematics, which makes multi-dimensional exploratory analysis easy. Creating multidimensional visual displays requires more effort, as everyone who does deep-uncertainty work. Figure 5 is one display that we have found comprehensible. It was constructed in *Excel* with *Analytica* output. It shows results by color or number of a cell, as a function of five parameters X1, X2, X3, X4, and X5.

Figure 5 An Illustrative Output of Exploratory Analysis

5 A Mixed-Methods Battery of Fusion Methods

We needed alternative fusion methods. We drew on the classic literature on basic information theory (Jaynes and Bretthorst, 2003), Bayesian analysis (Gelman and Shalizi 2010), Dempster-Shafer theory (Shafer 1976), and Dezert Smarandache theory (Smarandache and Dezert 2009a; Smarandache and Dezert 2009b), Bayesian networks (Pearl, 2009), and several others as surveyed in our earlier work (Davis, Perry, et al., 2013). We also drew on the literature on expert forecasts and judgment (Clement and Winkler, 2007).

We considered four types of fusion method: (a) purely subjective, (b) nonlinear algebraic, (c) quasi-Bayesian, and (d) a new entropy maximizing method (MEMP). All of these required adaptation or new work. Our nonlinear algebraic methods use the TLWS and PF methods discussed earlier. Our quasi-Bayesian method is “quasi” because we used heuristic methods to determine the weight given to different evidence and our model does not represent the full set of relationships and likelihoods.¹ We also constructed alternative “generic” likelihood functions and routinely show results for all of them because the “real” likelihood function is often unknowable. Finally, the ME/MP method comes from the perspective of information-theory entropy and uses methods from the machine-learning literature, such as those on regularization (see Davis, et al.

¹ Bayesian-network approaches can use more complex relationships. These require uncertain parameterization of likelihood functions and management of the resulting uncertainty. Potential problems include false precision and systematic model errors (Small and Fishbeck, 1999), which can be much worse than the problems from our simplifications. Using such models correctly can also be computationally intractable or at least very difficult.

forthcoming for citations). Technically, the approach uses nonlinear programming for fusion. It maximizes an objective function that includes a weighted sum of entropy-maximization terms and terms minimizing contradictions with reports, such as a claim that a person’s motivation is in the medium-to-high range, but no lower or higher. The method yields estimates of threat level that are as conservative (i.e., uncertain, in an information-theoretic sense) as possible given what has been reported, but with recognition that the reports’ assertions may not be correct.

Table 3 summarizes the methods. Since our information fusion involves multiple steps, we distinguish between (1) combining factors to estimate threat T and (2) fusing across reports, whether to improve the estimate of factors M , L , CO , and A , or to fuse the reports’ separate estimates of T .

Table 3 Mixed Fusion Methods

Combine factors to generate T for a given report	Fuse threat estimates across reports	Fuse factors across reports	Combine refined factors to generate T
Linear weighted sums (LWS)	Linear weighted sums (LWS)	Linear weighted sums (LWS)	Linear weighted sums (LWS)
Thresholded linear weighted sums (TLWS)	Thresholded linear weighted sums (TLWS)	Thresholded linear weighted sums (TLWS)	Thresholded linear weighted sums (TLWS)
Primary factors (PF)	Primary factors (PF)	Primary factors (PF)	Primary factors (PF)
Maximum entropy/minimum penalty (MEMP)	Maximum entropy/minimum penalty (MEMP)	Maximum entropy/minimum penalty (MEMP)	Maximum entropy/minimum penalty (MEMP)
	Quasi-Bayes	Quasi-Bayes	

Note: Although methods are mostly the same for all columns, inputs and outputs are context dependent.

6 Data

The results of even the most sophisticated analysis are limited by the data itself. We were doing research, not developing an application program. What should we use for data? It would have been a major effort to obtain “real” data and, had we done so, it might well have been classified or trivialized. More important scientifically, the data would not have really tested our platform adequately because the data would be whatever happened to be available. We instead created synthetic data to pose numerous challenges for our methods and platform—challenges that we had identified early, such as dealing with ambiguities in information (e.g., bimodal inputs

corresponding to equivocation), contradictions, even possible deception. We also designed cases that we suspected would show differences among methods.

This synthetic data took the form of narrative vignettes, followed by synthetic versions of analyst interpretations of raw data.

Here we give merely one short example, a single report from a human agent, Agent B, about an individual, Harry Smith, who is under scrutiny.

Agent B

Harry talks big and seems to see no problem at all with violence, even terrorist violence, in support of the group's objectives. I have to take that seriously, although he might be just "talking big."

It's hard to tell how motivated Harry is to participate in the Slammers' activity and from my position in the club membership I do not know if he would be provided with means and tools to participate in a violent act if he should choose to do so. Interestingly, I did hear comments that suggest Harry is nonetheless dubious about paying the cost and taking the risks associated with violent action—even though he has endorsed violence as a legitimate route to right society's wrongs.

From my vantage point, Harry does not appear to be a threat—he likes to talk big about the need for bringing down the system, but never wants to actually do anything other than talk. The few times he's been pushed he always comes up with some excuse about taking care of some pressing need.

Such narrative information must be "interpreted" in terms of probability distributions for Harry's motivation, sense of legitimacy, capability, and willingness to accept costs and risks (M , L , CO , and A). Those become the inputs for the analysis.

7 Designing and Implementing a Platform for Exploratory Analysis

Implementing the concept of Figure 1 involved more detailed design and some advanced programming. The initial version of the platform was much more procedural and straightforward—something an analyst who was not a programming expert could understand within days or a few weeks. That, however, hard-wired many of the choices we saw as important. Thus, we redesigned to achieve the flexibilities discussed earlier in connection with Figure 1. This involved tradeoffs. To our biased eyes, the design seems elegant, reviewable, and maintainable while maintaining the uncertainty information. That is so, however, only for those who have paid the price of a steeper learning curve than is required for more usual programming.

As an example, a significant meta-method issue is whether to combine the probability distributions for factors M , L , CO , and A for each report to generate a threat estimate T , and then fuse those T estimates across reports, or to fuse factors across reports to improve the estimated distributions for M , L , CO , and A , and then combine to estimate T . A number of such choices exist and—lacking settled theory and solid data—no *a priori* reason exists for believing that one

is “right.” The best choice depends on the individual case and its data. This is analogous to empirical statistical analysis where data-cleaning and model choice are often justified heuristically on a case-by-case basis.

One way to think about the structural issues around which we had to design is to think of the meta-level fusion as performed by operators as indicated in Figure 6. We have to operate on the data to map the raw data into the inputs of the platform; we have to assign the data to each of the streams; we have to decide on what order we process reports; we have to decide on when and how to combine and fuse; and we have to decide on when to fuse across reports rather than streams. Some of these operators don’t commute. The answers are different if we combine first rather than fuse first, etc.

Figure 6. Operators Employed in the Information Fusion

I : interpret data, specify parameters, via interface model
 S : assign data to stream(s)
 O : set report order
 C : combine factors to estimate threat
 F : fuse across reports
 F_{St} : fuse across streams
 $CF \neq FC$
 $F F_{St} \neq F_{St} F$
 ...

The resulting dimensionality is suggested by Table 5, which shows choices available for four of the structural issues and Table 5, which discusses data-grouping issues. The structural issues in Table 4 establish 48 analytical paths, each with uncertain parameters, as well as the choices indicated in Table 5. Exploratory analysis can generate tens of thousands of distinguishable cases, or more. It’s best not to dwell on the number because it isn’t conceptually important, or even computationally limiting at this stage, although applications of exploratory analysis benefit from supercomputers.

Table 5: The Primary Dimensions of Structural Uncertainty

Stream and Model	Combine or Fuse First?	Combining Method?	Fusion Method
A (with Propensity for Terrorism model, PFT)	Combine first	TLWS (Thresholded linear weighted sums)	TLWS (Thresholded linear weighted sums)*
B (with a variant model)	Fuse first	PF (Primary Factors)*	PF (Primary Factors)*
		MEMP (Maximum Entropy/Minimum	MEMP (Maximum Entropy/Minimum

Penalty)*	Penalty)*
	Quasi-Bayesian*

Note: All the methods include tuning parameters such as weights, thresholds, or maximum adjustments.

Table 6 Illustrative Data Groupings

Individual	Order of Report Processing	Report Weightings (for credibility...)
Harry	As received (chronological)	Uniform
Ahmed	Reverse order (newest first)	Through Lens 1
	Arbitrary 1 (e.g., best source first)	Through Lens 2
	Arbitrary 2	...

To elaborate slightly, we dealt with such organizational complexity as follows. At the visual-modeling level we treated the stream and the combine/fuse order as explicitly distinct flows, as in the architecture of Figure 1. Where possible, we represented the various operators mentioned above with functions. These functions have parameters to deal with structural issues and some data-group issues. Analytica deals with all of this using array mathematics. The result is that the mathematics of the heterogeneous fusion can be expressed compactly. For example, the node at which threat estimates are to be fused across reports is defined as a table, as in Table 7. The definition specifies that all four fusion methods will be used, each with its own function. The arguments specify on what data the functions are to operate. Thus, what appears to be a simple table of functions is actually a table of functions with multidimensional arrays of probabilities as inputs and new multidimensional arrays of probabilities as outputs. The design allows more extensive and comprehensible exploratory analysis that explains what assumptions and choices are affecting results.

Table 7. The Definition of the Node Accomplishing Fusion Across Reports

Fusion Method	
PF	Primary_Factor_Funct (Combined_Threat, R_ind, Tau, Final_Rpt_Wts_Agg, Min_Q_for_PF)
TLWS	TLWS_Function(MEMP_Fix_R_Threat, Final_Rpt_Wts_Agg, R_Ind, Threshold:4)
Bayes	Bayes_fx_TD(MEMP_Fix_R_Threat, R_Ind, Final_Rpt_Wts_Agg, Likelihood, Report_Order)
MEMP	Continuous_MEM_Threat[Harry_or_Ahmed_Index=Choose_Harry_or_Ahmed]


8 Illustrative Results and Conclusions

8.1 Results

Mostly to illustrate that the design was achieved and that the platform runs, Figure 7 shows illustrative but simplified results for one case. It shows the probabilities, after fusion, for assigning an individual, Harry Smith, to one of three categories: threat, gray-area, or non-threat (corresponding to cumulative probabilities of T between 6 and 10, 4 to 6, and 0 to 4). Results are shown as a function of fusion method. The dashed lines indicate what the assessment was with just the initial information. The “slicer bars” at the top indicate major contributors to the result beyond the choice of fusion method. For the example, in this particular projection of results, the TLWS method was used to estimate threat by combining factor values, Quasi-Bayesian fusion was accomplished using a postulated “asymmetric” likelihood function, and the reports were processed in the order received. Such parameter values can be changed interactively by clicking through their menus (note arrows). This is interactive exploratory analysis.

We see that in this fusion greatly increases the likelihood ascribed to Harry being a threat. The primary factors method is extreme in this regard, as one would expect, but a rather striking result is that the Bayesian method *drops* the estimate by a factor of 4 (0.28 to 0.07)! This doesn’t exonerate Harry, but it strongly suggests that he is not a terrorist. But what if we used other fusion methods? Would we get the same answer? Not necessarily. Seeing discrepancies tells us to look more deeply into the particulars of Harry’s case, perhaps all the way back to raw data. After doing so, we will probably be able to focus on one or two fusion method and refine the tuning. Again, however, this envisions an intelligent analyst-centric process, not an automated statistical analysis.

Figure 7: Illustrative Results from Prototype Experiments (not chosen for drama)



8.2 Conclusions

By the end of our project, the methods were falling into place, the prototype analytical platform was operating and we saw significant consequences of going about heterogeneous fusion in different ways. This was “good,” not “bad,” because we had used the mixed-method approach precisely because we expected discrepancies and the need to look at the problem from different perspectives. Follow-up analysis would be different depending on context. For example, if we were desperately trying to find the most plausible suspect among a set of people, we might “look for trouble,” choosing methods and tuning parameters accordingly. If instead we were trying to objectively and dispassionately assess threat likelihoods, we would do something else. And, finally, if we were second-guessing the decision to regard someone as a probable terrorist, we would look to find what formation elements were most critical in that assessment. If one element had particularly high leverage, then we would scrutinize it in more detail and seek new information to corroborate or disconfirm it. The result might be exoneration.

We conclude that such fusion methods have promise for increasing the probability of detecting the rare potential terrorist, decreasing false alarms, and increasing the probability of exonerating individuals who might otherwise be falsely assessed. Much future work will be needed, however, to determine *actual* value in real-world settings. We were merely establishing some groundwork for next steps.

Acknowledgments

This paper is based on prior research sponsored by the Office of Naval Research and the Office of the Secretary of Defense.

References

- Clemen, Robert T., and Robert L. Winkler (2007) "Aggregation of Expert Probability Judgments," in Wade Edwards, Miles, Ralph F, and Detlof von Winterfeldt, eds., *Advances in Decision Analysis: from Foundations to Applications*, pp. 154-176.
- Davis, Paul K. (2012), *Some Lessons From RAND's Work on Planning Under Uncertainty for National Security*, Santa Monica Calif.: RAND Corp.
- Davis, Paul K. (2003), "Exploratory Analysis and Implications for Modeling," in *New Challenges, New Tools*, edited by Stuart Johnson, Martin Libicki, and Gregory Treverton, Santa Monica, Calif.: RAND Corp., 255-83.
- Davis, Paul K., and Kim Cragin, eds. (2009), *Social Science for Counterterrorism: Putting the Pieces Together*, Santa Monica, Calif.: RAND Corp.
- Davis, Paul K. Eric Larson, et al. (2012), *Understanding and Influencing Public Support for Insurgency and Terrorism*, Santa Monica, Calif.: RAND Corp.
- Davis, Paul K., and Angela O'Mahony (2013). *A Computational Model of Public Support for Insurgency and Terrorism: A Prototype for More General Social-Science Modeling.* Santa Monica, Calif.: RAND Corp.
- Davis, Paul K., Walter S. Perry, Ryan Andrew Brown, Douglas Yeung, Parisa Roshan, and Phoenix Voorhies (2013). *Using Behavioral Indicators to Help Detect Potential Violent Acts*, Santa Monica, Calif.: RAND Corp.
- Davis, Paul K., Walter L. Perry, John Hollywood, and David Manheim (2015a), *Uncertainty Sensitive Heterogeneous Information Fusion: Assessing Threat with Soft, Uncertain, and Conflicting Evidence*, Santa Monica, Calif.: RAND, forthcoming.
- Davis, Paul K., Walter L. Perry, John Hollywood, and David Manheim (2015b), "Using Causal Models in Heterogeneous Information Fusion To Detect Terrorists," in *Proceedings of the 2015 Winter Simulation Conference*, edited by L. Yilmaz, K.V. Chan, I. Moon, T.M. Roeder, C. Macal, and M. D. Rossetti.

- Gelman, Andrew, and Cosma Rohilla Shalizi (2010). “Philosophy and the Practice of Bayesian Statistics,” *British Journal of Mathematical Statistics and Psychology*, 66, pp. 8-38.
- Jaynes, Edwin T., and G. Larry Bretthorst (ed.) (2003) , *Probability Theory: The Logic of Science*, Cambridge: Cambridge University Press.
- Kahneman, Daniel. (2011). *Thinking, Fast and Slow*, New York: Farrar, Straus and Giroux.
- Lempert, Robert J., David G. Groves, Steven W. Popper, and Steven C. Bankes (2006), “A General Analytic Method for Generating Robust Strategies and Narrative Scenarios,” *Management Science* 4, April 514–28.
- Lempert, Robert J., Steven W. Popper, and Steven C. Bankes, *Shaping the Next One Hundred Years: New Methods for Quantitative Long-term Policy Analysis*, Santa Monica, Calif.: RAND.
- Mitchell J. Small & Paul S. Fischbeck (1999), “False Precision in Bayesian Updating with Incomplete Models,” *Human and Ecological Risk Assessment, An International Journal*, 5:2, 291-304.
- National Research Council (2014), *U.S. Air Force Strategic Deterrence Analytic Capabilities: An Assessment of Methods, Tools, and Approaches for the 21st Century Security Environment*, Washington, D.C.: National Academies Press.
- Pearl, Judea (2009), *Causality: Models, Reasoning, and Inference*, Cambridge, Mass.: Cambridge University Press.
- Rosenhead, Jonathan, and John Mingers (2002). “A New Paradigm of Analysis,” In *Rational Analysis or a Problematic World Revisited: Problem Structuring Methods for Complexity, Uncertainty and Conflict*, edited by Jonathan Rosenhead, and John Mingers, 1–19. Chichester, UK: John Wiley & Sons, Inc.
- Shafer, Glenn (1976). *A Mathematical Theory of Evidence*, Princeton, New Jersey: Princeton University Press.
- Smarandache, Florentin, and Jean Dezert, eds. (2009a), *Advances and Applications of DsMT for Information Fusion*. Rehoboth: American Research Press.
- Walker, Warren E., Robert J. Lempert, and Han H. Kwakkel (2013), “Deep Uncertainty,” in S. Gass, and M. Fu, eds., *Encyclopedia of Operations Research and Management Science*, 3d ed., Springer.

Author Biographies

Paul K. Davis is a senior principal researcher at RAND and a professor of policy analysis in the Pardee RAND Graduate School. He is a graduate of the University of Michigan (B.S.) and Massachusetts Institute for Technology (Ph.D. in Chemical Physics). His research has included such diverse subjects as strategic planning; deterrence theory; counterterrorism theory; modeling, including cognitive modeling of adversaries and multi-resolution modeling more generally; and complex information fusion to assist threat detection. His email address is paul_k_davis@me.com.

David Manheim is a doctoral fellow in the Pardee RAND Graduate School and an Assistant Policy Analyst at RAND. His focus is decision support, risk analysis, and value of information. He holds a B.S. in Mathematics from Lander College, where he concentrated on abstract mathematics and financial modeling. His research includes understanding resilience in different

contexts, and several projects on how mathematical and computer modeling can inform decisions under uncertainty. His email address is dmanheim@rand.org.

Walter L. Perry is a senior information scientist at RAND. He received his Ph.D. at George Mason University after retiring from the U.S. Army's Signal Corps. He has taught electrical engineering, computer science, statistics, and mathematics. His research has included leading official reviews for the Army of operations in Kosovo, Afghanistan, and Iraq. More technically, his research has included methods for data fusion and information-processing, and for complex information fusion to assist threat detection. His email address is walt@rand.org.

John S. Hollywood is a full operations researcher at RAND and a professor of policy analysis at the Pardee RAND Graduate School, where he applies qualitative and quantitative analytics to security policy, including criminal justice, homeland security, counterinsurgency, and defense systems. He holds an SB in Applied Mathematics and a Ph.D. in Operations Research from the Massachusetts Institute for Technology. His email address is johnsh@rand.org.